System Memory Power and Thermal Management Techniques in Mobile Platforms

Jayesh Iyer

Corinne Hall

Yuchen Huang

Jerry Shi

Table of Contents

(Click on page number to jump to sections)

System Memory Power and Thermal Management Techniques in Mobile Platforms

Overview: New Throttling Techniques Increase Performance

As the form factors of computing platforms get smaller and component densities get higher driven by demand and technology, system power consumption and thermal issues have become more challenging than ever before, especially in the case of laptops, slim desktops (shown in **Figure 1**), and blade servers. The high performance of the Intel® Core™ Duo processor and Intel® 945GM Express chipset family in the platform built on Intel® Centrino® Duo mobile technology demands high electrical power and generates substantial heat.

Naturally, each system has a total cooling capacity per its specific design, and each major component has a cooling limit or power consumption allowance for balancing the system performance in that design. When the total system cooling capacity is smaller than the sum of each component's thermal design power (TDP), as seen in many small-form-factor systems, all the components simply cannot simultaneously operate at their TDP level since the system will fail to cool them. In such a system, it's also unacceptable to allow one component to free-run without power consumption and thermal restrictions, leaving other components to suffer in performance.

Balanced cooling limits for major components should be achieved at the design level. Having briefly laid the background on cooling limits, this article will now focus on system memory (specifically in laptops), which is a major power-consuming component on platforms—its cooling limits, power, power prediction, and bandwidth recovery using new throttling techniques.

Figure 1. *Popular laptop and slim desktop form factors.*

To better understand the variation in memory cooling limits from system to system, a glimpse at three popular categories of laptop is helpful—the thin-and-light laptops (with a Z-height of around 1.1–1.2 inches and a memory cooling limit of 4–5 watts), the mini-note PCs (with a lower Z-height of around 0.9 inches, a smaller cooling fan and a memory cooling limit of 2–2.5 watts) and the subnote PCs (which don't have a cooling fan and have a cooling limit around 1 watt).

The Need for System Memory

Figure 2 shows a high-level view of the main components on a mobile platform built on Intel Centrino Duo mobile technology. The system memory gets accessed in almost all activities taking place in the platform. All data transfers to and from the system memory are managed by the Intel 945GM Express chipset. Thus, if the chipset/system memory is idle, the platform itself is generally in an idle state. But if there is activity on the chipset/system memory, the platform is consuming more power, causing the chipset and memory components to heat up.

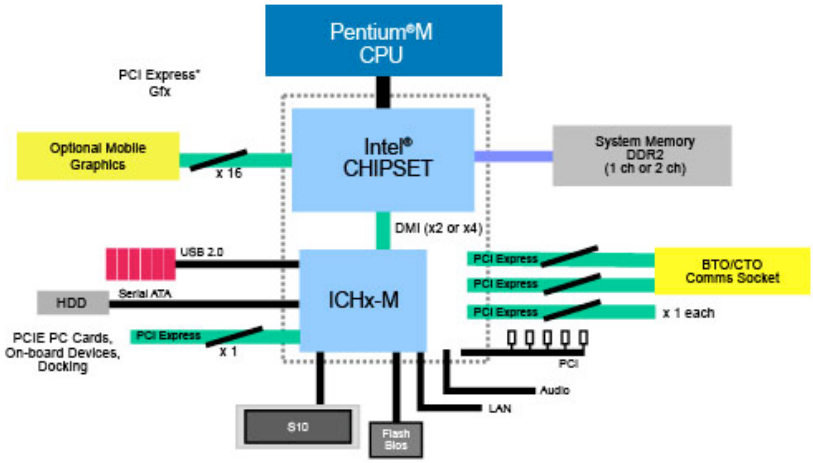

Figure 2. *Schematic showing the main components in current-generation laptops.*

As mentioned earlier, mobile platforms have limited cooling capabilities. In the past, memory speeds and capacities generally allowed the system memory subsection to stay within the cooling limits of the platform, and the DRAM case temperatures were below the maximum operating specifications of 85°C. But with the increase in memory capacity and speed, we are now reaching the point where memory thermals are starting to exceed the cooling capabilities of mobile systems. When the cooling budget is exceeded, that means the system is no longer able to cool the memory subsection, and DRAM case temperatures begin to exceed their max-case temperature specification.

In a recent laboratory study, several different thin–and-light laptop book designs were tested with various memory speeds, capacities, densities and vendors, using the OpenGL* Benchmark software. This software has high bandwidth utilization (about 52 percent of theoretical max) and causes DRAM devices to draw constant high power. **Figure 3** shows the results of a small sample of this thermal study.

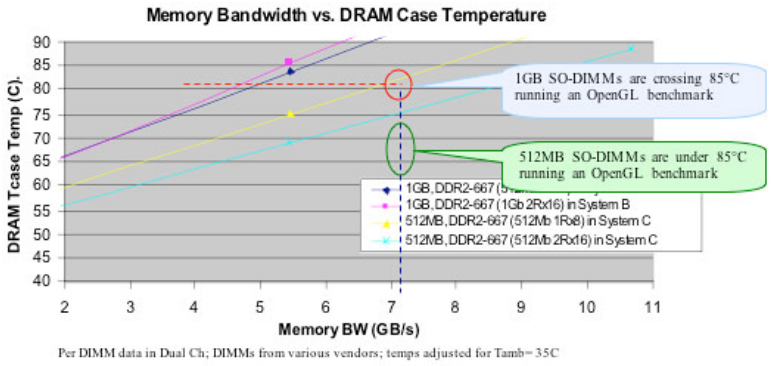

Figure 3. *Lab data showing memory bandwidth and DRAM case temperatures. Temperature and bandwidth data is dependent on system, memory and software configuration.*

The results of the above study show that typical 512MB SO-DIMMs in various thin-and-light laptop designs are well below 85°C and do not appear to be in jeopardy of exceeding their thermal limits. In some cases, however, 1-GB-capacity SO-DIMMs are reaching, and sometimes exceeding, their maximum-case temperature specification of 85°C. These results show that memory modules are already operating near their maximum specifications, and as memory power and thermals increase with system capacity and speed, memory modules will begin exceeding their maximum specification with multiple realistic workloads. There is clearly a need for a robust thermal management solution.
If left unchecked, DRAM devices will start running above their maximum operating case temperatures, and memory-related reliability issues will begin to crop up. Thus the memory bus needs to be throttled to ensure that the DRAM devices operate within their thermal limits. Memory throttling provides a solution to cool the DRAM devices by reducing memory traffic allowed on the memory bus, thereby reducing the power consumed by the DRAM devices and thus reducing thermal output. The article will now discuss two such memory-throttling techniques.

Delta Temperature in Serial Presence Detect

Delta temperature in serial presence detect (DT in SPD) is a throttling technique that gives a huge performance benefit in platforms that do not have a physical thermal sensor on the memory module. Before going into details of DT in SPD, let's briefly look at the throttling methodology followed in the chipsets of previous-generation platforms.
A double data rate (DDRx)–based memory subsystem, as a major power-consuming component on a platform, has a peculiarity in that its power consumption is difficult to predict. The reason for this is that memory vendors have their own unique processes and design technologies. This results in large variations among different DRAM vendors in power consumption given a fixed device density and speed for each particular DDRx technology.[1]

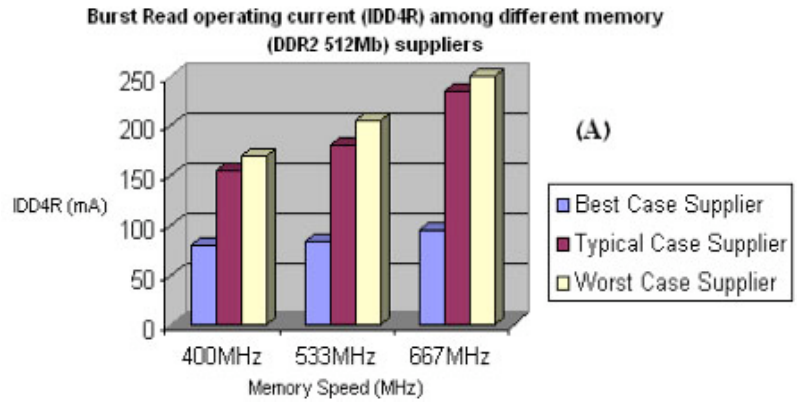

***Figure 4**. Shows the differences in the Burst Read current (IDD4R) of the best-, typical-, and worst-case memory supplier (DDR2 512-MB DRAM) across three different memory speeds. As the memory device power is a function of its IDD numbers, a huge difference in the IDD numbers between the best-case and worst-case memory suppliers translates into a huge difference in their power consumption.*

In the previous generation of mobile platforms, the system could tell the memory type (device density, number of ranks, device width) and its vendor by retrieving the data from a small EEPROM (SPD) on each memory module during boot up. But the information never indicated how much power the module consumed.

The chipsets in these platforms implemented a throttling methodology without any feedback on power or temperature from the memory module. Without this knowledge of power consumption from the memory module, the system was forced to assume the worst-case power consumption for design safety. Thus, memory thermal limits (throttling threshold value) were set based on the worst-case supplier's power data. This approach thus resulted in over-guardbanding[2] and over-memory-throttling. The memory performance degrades as low-power memory modules are treated as worst-case modules.

Now the inevitable question is, how can a system, especially when it does not have a physical thermal sensor on the memory modules, accurately and cost-effectively tell memory power consumption to avoid over-throttling memory? That is, how can it reduce the guardband? Delta temperature in SPD was the first step to the solution that addresses this issue.

Overview of DT in SPD

Delta temperature in serial presence detect (DT in SPD) is a power/thermal prediction scheme providing DRAM power/thermal data in the SPD on the module. It stores key temperature rise data and DRAM maximum-case temperature limit data (Tcasemax) in the SPD on the memory module. The system uses this information to better estimate the temperature of the DRAM devices and to determine when throttling is necessary. When process shrinks or other power optimizations occur and DRAM power dissipation decreases, the system uses the information stored in SPD to reduce memory throttling and regain system performance.

Usage Models

Memory module vendors report the delta temperature rise parameters and Tcasemax in SPD. These parameters are read by the BIOS from the SPD at boot time. Subsequently, the system adjusts memory throttle limits based on these parameters.

Performance Benefits of DT in SPD

DT in SPD greatly reduces the guard band and helps recover the bandwidth, as compared to the throttling methodologies implemented in chipsets in the previous platforms, which set memory thermal limits (throttling threshold value) based on the worst-case supplier's power data. A lab study was done on thin-and-light laptop designs using memory modules of different configurations from different vendors. As mentioned earlier, there is a huge difference in the power numbers between different vendors (for the same configuration). The best-case, typical-case, and worst-case DRAM vendors (for each configuration of the memory) in terms of power numbers were used for the study. A Windows*-based chipset stress utility, which generates very high traffic (almost all page hits) on the memory bus, was used to exercise the system memory.

Figure 5 and **Figure 6** each give a comparison of performance achieved running the utility with 30 percent Writes and 70 percent Reads on the thin-and-light design. It shows the sustained bandwidth across different memory throttle limits using the two throttling control methodologies: 1) With DT in SPD; 2) Without DT in SPD (using worst-case supplier power data).

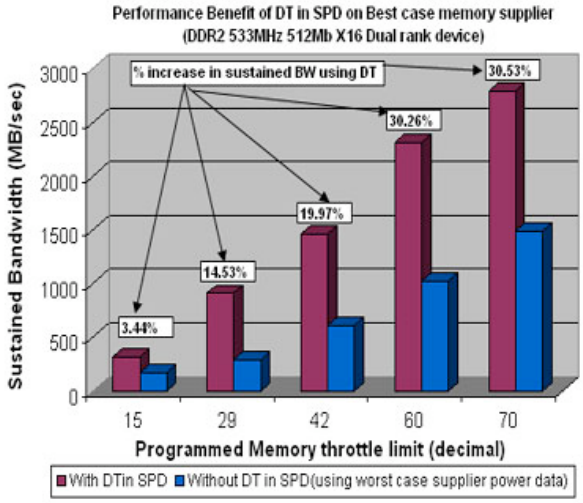

Figure 5. *Sustained bandwidth across memory throttle limits (best-case supplier).*

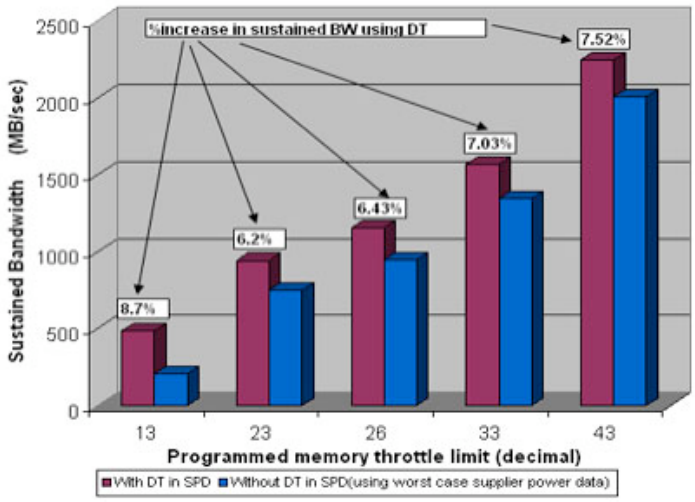

Figure 6. *Sustained bandwidth across memory throttle limits (typical-case supplier).*

The results in **Figure 5** show that a typical-case memory supplier gives a performance benefit of 8 percent (in terms of sustained bandwidth) using DT in SPD. But for measurements done on a best-case memory supplier, the performance benefit goes up to 30 percent (in terms of sustained bandwidth) using DT in SPD as shown in **Figure 6**. Thus the results clearly demonstrate that using DT in SPD greatly reduces the guard band and helps recover the bandwidth. This enhances memory performance, especially for very-low-power memory modules (best-case supplier).

Though DT in SPD has its benefits in systems that don't have a physical thermal sensor on the memory modules, it is still an open-loop throttling mechanism as it does not actually know the operating temperature of the DIMM. Therefore it must always assume the laptop is operating under the maximum-allowed room ambient temperature, which is typically considered to be 35°C for mobile environments. So if a laptop is actually running in an environment where the room ambient is only 20°C, then DT in SPD assumes it is running at 35°C and will initiate throttling ~15°C sooner than it needs to, thereby losing potential performance that could have been obtained had the system not initiated throttling so soon.

Thermal Sensor on DIMM

The follow-on throttling mechanism to DT in SPD is thermal sensor on dual in-line memory module (TS on DIMM). TS on DIMM offers a more advanced approach to system thermal management through the use of a physical thermal sensor on the memory module to further reduce the guardband present in the existing methods. The reduction of guardband has a direct and positive impact on system performance.

Overview of TS on DIMM

Instead of using power prediction to estimate the temperature of the DRAM case, TS on DIMM uses a physical thermal sensor integrated on the DIMM module to monitor the temperature of the DIMM. If the thermal sensor detects that the DIMM temperature is exceeding a programmable critical trip point, then the thermal sensor triggers an event signal that tells the chipset to throttle the memory traffic, reducing the DRAM case temperature.

The addition of a physical thermal sensor results in a closed-loop throttling methodology that allows for real-time throttling based on measured temperature. This closed-loop methodology leads to a reduced guardband primarily due to the ability to sense ambient temperature. Most of the guardband present in existing throttling mechanisms is due to their open loop policy, which means that since they do not receive feedback on the actual ambient or DIMM temperatures, they must always assume the worst-case scenario. TS on DIMM provides the needed temperature feedback and allows the system to hold off on throttling until the actual DIMM temperature reaches a programmedcritical temperature trip point. The closed-loop policy offered by TS on DIMM is what allows it to reduce the guardband over existing memory throttling methods.

Performance Benefits

The reduction in guardband that TS on DIMM offers over existing methods gives it a performance edge in both bandwidth as well as benchmark scores when running high-bandwidth applications. When using an OpenGL benchmark on 1-GB DDR2-667MB SO-DIMMs in dual-channel mode, lab data showed that for every degree C of guardband removed, the system saw up to a 220-megabyte-per-second (MBps) bandwidth improvement per degree C, as shown in **Figure 7**. Since TS on DIMM typically saves ~15°C of guardband over existing methods when running a high-bandwidth application at room temperature, and at 220 MBps per degree C, the total bandwidth recovered with TS on DIMM can be approximately 3 GBps for this particular application, which is 30 percent of the theoretical maximum bandwidth for this system and memory configuration. This bandwidth recovery assumes that the workload running on the system could utilize more bandwidth if allowed. The bandwidth recovered is also dependent on the application used and the system configuration.

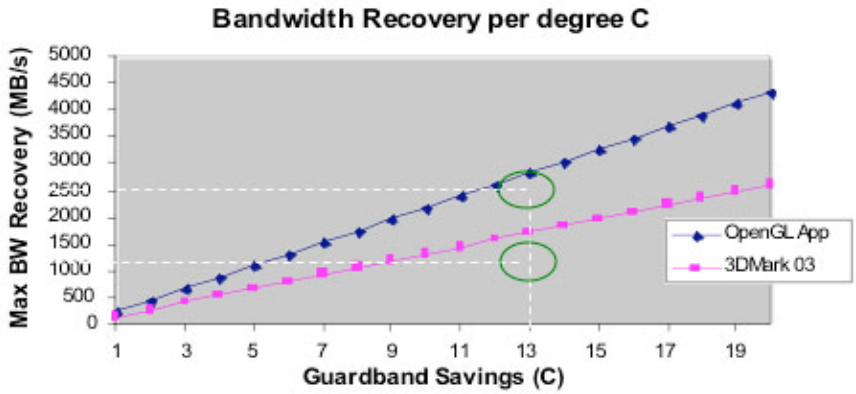

Figure 7. *Bandwidth recovered per degree of guardband removed.*
Temperature and bandwidth data is dependent on system, memory, and software configuration.

Due to the increased bandwidth per degree C of guard-band savings, benchmark scores also show improvements with TS on DIMM, as shown in **Figure 8**. Using the same 1-GB, DDR2-667 SO-DIMMs in dual-channel mode as described above, for every degree C of guardband removed, the system saw ~3.3 frames per second per degree C for a particular OpenGL application, and up to 13 3DMark* score recovery with 3DMark03 per degree C. Again, since TS on DIMM typically saves ~15°C over existing methods, that is an approximate 50-frame-per-second improvement as well as about a 200-3DMark score improvement with these applications. Results will vary with different applications and different system configurations.

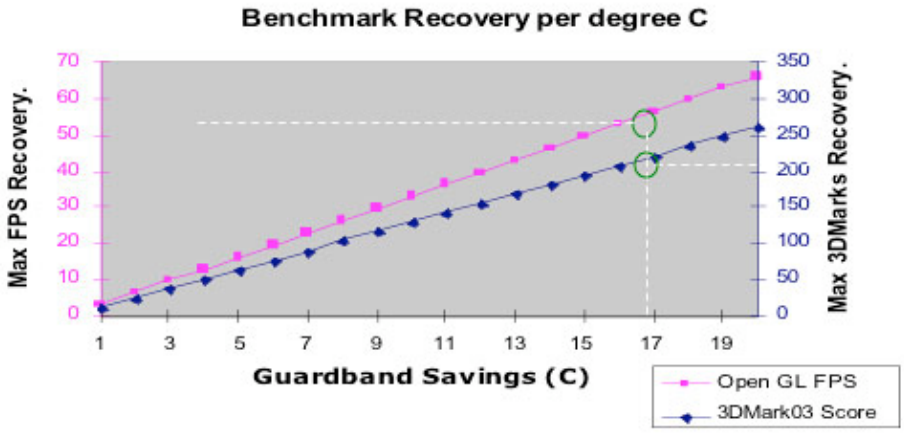

Figure 8. *Benchmark scores recovered per degree of guardband removed.*
Benchmark score recovery is dependent on system, memory, and software configuration.

To further understand how the reduction in guardband impacts bandwidth availability, let's look at the thermal study from earlier in this article and determine where TS on DIMM would throttle the worst-case 1-GB SO-DIMM module in comparison to DT in SPD. Please see **Figure 9** below.

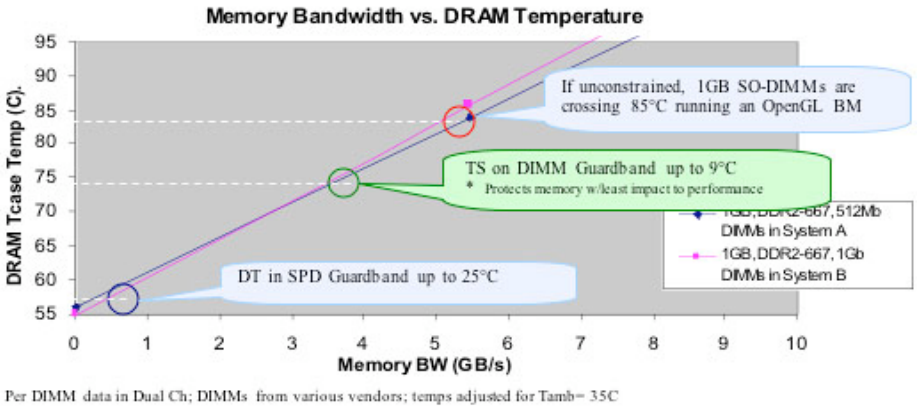

Figure 9. *Guardband comparison between TS on DIMM and DT in SPD.*
Temperature and bandwidth data is dependent on system, memory, and software configuration. Guardband numbers shown are worst-case.

The above chart shows that the reduction in guardband directly impacts bandwidth availability as the temperature of the DIMM reaches and exceeds 60°C.

Other Benefits of TS on DIMM

Aside from the performance benefits mentioned above, TS on DIMM offers other benefits over existing methods. First, since the thermal sensor is measuring true operating temperature, it is able to protect the DRAM cases from exceeding their maximum specifications even if an outside source is heating the DIMMs. For example, if a fan breaks or if the laptop is sitting in direct sunlight, the thermal sensor is always monitoring the DRAM temperatures, and if they exceed the critical trip point, the chipset will throttle the system and protect the DRAMs.

With the existing methods of temperature prediction based on memory traffic alone, the chipset has no idea if some abnormality is causing the DRAM temperatures to rise beyond their specifications. Since TS on DIMM is monitoring temperature of the memory, it also provides a way for controlling laptop skin temperatures. Keeping memory within certain thermal limits will also help keep skin temperatures within certain thermal limits. Another benefit that TS on DIMM offers over existing methods is a reduced number of SPD byte requirements. DT in SPD requires 15–20 bytes for implementation; TS on DIMM reduces that byte requirement substantially to about 4 bytes. Plus, the SPD bytes for TS on DIMM are easier to calculate and easier to manage over DT in SPD, and the bytes are less dependent on DRAM changes.

Summary

Specific form-factor requirements of laptops limit them from having exhaustive cooling solutions. As a result, individual components inside the laptop (like system memory devices) can heat up to the point of exceeding their operating specifications. To ensure that the memory devices operate within their thermal limits, their case temperature needs to be monitored and memory accesses throttled if any memory devices approach their thermal limits. Current lab data taken on multiple laptops and by analyzing multiple SO-DIMMs show that 1-GB capacity SO-DIMMs and greater are nearing their maximum specified case temperature when running a realistic workload at an ambient temperature of 35°C. This data was taken from thin-and-light laptop designs, but the thermal issues for small-form-factor designs are of even greater concern due to their limited cooling abilities.

Current throttling mechanisms, including DT in SPD, throttle memory based on a thermal prediction scheme that analyzes memory traffic to estimate the temperature of the DRAM case temperatures. These open-loop throttling mechanisms lead to over-guardbanding, since they always have to assume worst-case environment conditions, like room ambient. Also, because these mechanisms are estimating thermal parameters based only on memory traffic, they cannot protect all DRAMs from exceeding their maximum case specifications of 85°C.

TS on DIMM is the follow-on throttling mechanism to DT in SPD, offering several benefits over existing methods. The closed-loop methodology of TS on DIMM allows it to reduce the guardband substantially over existing methods, increasing system performance by allowing the system to run unconstrained longer before initiating throttling. TS on DIMM also provides a way to control laptop skin temperatures and provides a safeguard mechanism to prevent the DRAMs from exceeding their maximum case specifications of 85°C.

Currently TS on DIMM is an optional feature for DDR2 SO-DIMMs. Intel is working closely with the memory industry to drive these features into DDR3.

More Info

Discover more at the Intel Web site:

- Networking & Communications
- Technology & Research

Author Bios

Jayesh Iyer
Hardware Design Engineer
Mobile Platforms Group
Intel Corporation

Jayesh Iyer joined the Intel Mobile Platforms Group in 2003. He is part of the Mobile Platform Architecture and Development team, designing customer reference boards for next-generation Intel® processors and chipsets. He also works on platform memory power/performance initiatives (like DT in SPD and TS on DIMM) and memory profiling. He has also worked as a research intern at the Indian Space Research Organization. Iyer received his B.E. degree in instrumentation and control engineering from Gujarat University and an M.S. in electrical engineering and computer engineering from Drexel University, Philadelphia.

Corinne Hall
Board Architecture Engineer
Mobile Platforms Group
Intel Corporation

Corinne Hall joined Intel in 2001 and has worked in the Desktop Products Organization as well as the Mobile Platform Group. Her current responsibilities include thermal sensor on DIMM enabling as well as other memory thermal and power-related activities for DDR2 and DDR3. Hall received a B.S. degree in computer engineering from Oregon State University.

Yuchen Huang
Staff Engineer
Technology & Manufacturing Group
Intel Corporation

Yuchen Huang is a staff platform memory power engineer in Intel's Technology & Manufacturing Group. She joined Intel in 1994 and has held various platform design positions there as Intel's Desktop Products Group platform memory power and performance initiatives owner, and as a platform interconnect designer. Currently, Huang leads a cross-divisional engineering team within Intel that addresses memory power, thermal and performance platform challenges. Huang received a B.S. degree in electrical engineering from Zhejiang University, China, and an M.S. degree in electrical engineering from Portland State University.

Jerry Shi
Staff Memory Architect
Mobile Platforms Group
Intel Corporation

Jerry Shi has been with Intel for 7 years, working as platform/memory architect. He designed the first Intel® networking processor-based cPCI system for demonstration at SuperComm. While at Intel, he also designed the first set-top box that could do picture-in-picture MPEG playback. He holds a number of patents related to memory architecture. Prior to joining Intel, he worked for Philips Semiconductors, Hyundai Electronics, and Oak Technology. Shi received a B.S. degree in computer science and engineering from Tsinghua University, China, and an M.S. degree in electrical engineering from the University of Massachusetts at Amherst.

Notes

[1] X is 1, 2, or 3 based on the double data rate (DDR, DDR2, DDR3) synchronous DRAM memory technology.

[2] Reduction in bandwidth to ensure power/thermal values stay within design safety limit.

—End of Technology@Intel Magazine Article—